Integer Programming Ensemble of Classifiers for Temporal Relations

Catherine Kerr[1], **Terri Hoare**[2], **Jakub Mareček**[3], **Paula Carroll**[4]

[1],[2],[4] University College Dublin

{catherine.kerr,terry.hoare}@ucdconnect.ie, paula.carroll@ucd.ie

[3] IBM Research – Ireland

jakub.marecek@ie.ibm.com

arXiv:1412.1866v2 [cs.CL] 18 Jun 2015

Abstract

Extraction of events and understanding related temporal expression among them is a major challenge in natural language processing. In longer texts, processing on sentence-by-sentence or expression-by-expression basis often fails, in part due to the disregard for the consistency of the processed data. We present an ensemble method, which reconciles the output of multiple classifiers for temporal expressions, subject to consistency constraints across the whole text. The use of integer programming to enforce the consistency constraints globally improves upon the best published results from the TempEval-3 Challenge considerably.

1 Introduction

Extraction of information from texts, such as the extraction and analysis of events and related temporal expression, is an increasingly important challenge in Natural Language Processing (NLP). Applications range from machine reading in military intelligence to compiling situational awareness reports in disaster management. A series of NLP competitions, Temporal information extraction challenge (TempEval) and Semantic Evaluation (SemEval), have helped to focus much of the research within event and temporal relation processing, not least by standardising data formats and performance measures.

TempEval Task C focuses on the detection and classification of temporal relations between events in news-feed documents (UzZaman et al., 2013). The 2013 TempEval Challenge was the first challenge to include all fourteen temporal relations of the TimeML temporal information standard markup language (UzZaman, 2012), which are, in turn, based on Allen's intervals algebras (Allen, 1983). While leading classifiers (Chambers, 2013; Bethard, 2013; Laokulrat et al., 2013) can detect the patterns required to predict the most likely relation between any two events, they do not check for global consistency.

This paper presents an Integer Programming (IP) ensemble, which reconciles the results of multiple temporal classifiers, subject to the full set of consistency constraints across Allen relations in each news-feed document. The application of the global rules of a relation algebra on the full set of relations provides a necessary and sufficient condition for consistency across the relations in a news-feed document (Ladkin, 1990).

For an illustration of the potential benefits of an ensemble, consider the three event-relationship graphs (ERGs) in Figures 1–3. (See Glavaš and Šnajder (2014) for a survey of event-relationship graphs.) In Figure 1, there is the hand-annotated BBC news-feed document. Figures 2 and 3 show the diversity in annotation between two of the machine-learning classifiers for the same feed. Especially in Figure 2, it is clear that the references to an "origin" often prevail. As has been observed (Florian et al., 2002), this diversity can generally be exploited by combining the classifiers to improve overall temporal awareness.

2 Related Work

There is a rich history of calculi for temporal reasoning, which underlies more recent work on processing temporal expressions in computational linguistics. Our work is based on the approach of Allen (1983), which defines the transitive composition on the set of relations. For example, for the transitive triangle formed by the events A; B; and C in an acyclic event-graph, if the event A is BEFORE the event B and the event B is BEFORE the event C then the composite relation between A and C must be BEFORE. Allen further defines the concept of temporal path consistency or temporal closure on event graphs. Ladkin (1990) defines that for path consistency to hold, the relation on any edge in the graph must be contained in the composition along any path beginning at the tail and ending at the head.

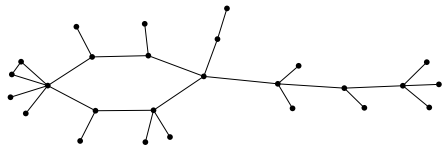

Figure 1: Platinum hand-annotated news feed (bbc_20130322_1600)

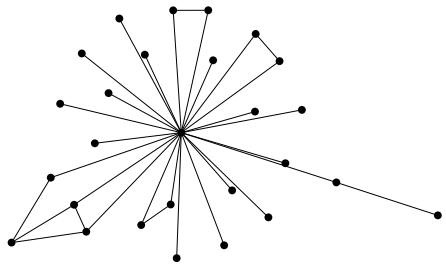

Figure 2: Cleartk-2 machine-annotated news feed (bbc_20130322_1600)

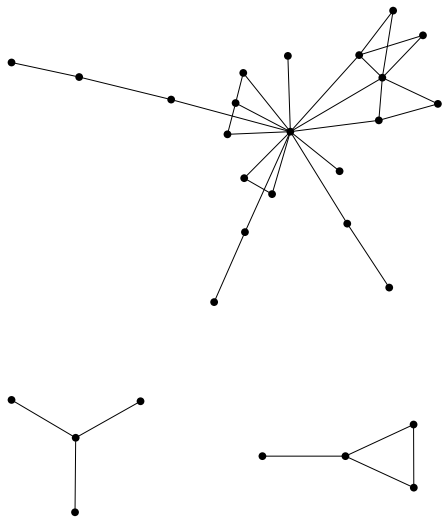

Figure 3: Navytime-1 machine-annotated news feed (bbc_20130322_1600)

In more recent work, TimeML has been adopted as the ISO standard markup language for temporal information. It captures temporal relations with Temporal Links (TLINKs), (Saurí et al., 2009). A TLINK can be represented as an arc on an event graph. A TLINK either anchors an Event (EVENT-ID) to a Time Expression (TIMEX) or orders the relationship between two Events (EVENT-IDs). TLINKs are based on the Allen mutually exclusive relations between two discrete event or time intervals, (Allen, 1983).

The most recent TempEval challenge included the full set of fourteen TimeML relations (UzZaman et al., 2013). Eleven participant classifiers used a variety of machine-learning techniques including support vector machines and methods inspired by maximum-entropy, yielding very different results. Naturally, one may want to combine the results of multiple classifiers in order to improve the performance, e.g., in terms of F1 score. A generic ensemble of classifiers that make predictions on pairwise relations from a set of fourteen possible relations can result in contradictory predictions of relations. One can, however, also reconcile them while enforcing transitivity rules of the underlying relation algebra.

In integer linear programming (IP), a linear function is optimised over an intersection of linear inequalities and equalities (polyhedron) and integrality constraints (integral grid). Although a polynomial in the number of variables and the number of constraints does not bound the run-time, unless P = NP, the computational performance turns out to be very good, in practice. The branch and bound and cut methods, implemented in modern IP solvers, significantly reduce the number of relaxations required to solve the IP problem. Hence, IP has been applied to a wide range of problems in logic, combinatorics, and most recently, computational linguistics. Examples include text summarisation, (Woodsend and Lapata, 2011), semantic role labelling, (Roth and Yih, 2004; Punyakanok et al., 2004), global reconciliation on temporal labels, (Chambers and Jurafsky, 2008), coupling of local event-event and event-time classifiers, (Do et al., 2012). In these cases, the IP formulation makes it possible to enforce global constraints over subgraphs of intervals only, for example UzZaman et al. (2013) enforces global constraints over three intervals. In our case, the rules of interval algebra were applied over the *whole* instance.

3 The Ensemble Method

The ensemble method takes as its input the output for each annotated TLINKs from each classifier. We use A to denote the set of annotated TLINKS, $\{1..15\}$ to denote the fourteen *reltypes* plus NONE (no relation), and C^* for the set of relationships that can be inferred from $a \circ b$ and $c \in C^* \subset \{1..14\}$. The binary indicator decision vari-

able is:

$$x_{i,j} = \begin{cases} 1, \text{ if relationship } j \text{ is assigned to TLINK } i \\ 0, \text{ otherwise.} \end{cases}$$

The integer program can be seen as a weighted assignment problem with additional constraints:

$$\max \sum_{i \in A} \sum_{j=1}^{15} \alpha_{i,j} x_{i,j} \quad (1)$$

$$\text{s.t.} \sum_{j=1}^{15} x_{i,j} = 1, \qquad \forall i \in A \quad (2)$$

$$x_{pq,a} + x_{qr,b} - \sum_{c \in C^*} x_{pr,c} \leq 1, \quad (3)$$

$$\forall pq, qr, pr \in A, a, b \in \{1..15\}.$$

In the objective (1), $\alpha_{i,j}$ is a weight, which throughout this paper is the sum of F1 scores of classifiers suggesting the class. Notice that the weights are calculated in the data preparation step, which allows for a wealth of non-linear functions to be explored. Contraint (2) expresses mutual exclusivity, which guarantees that only one relation type can be assigned to each arc or TLINK. Constraint (3) performs a transitive composition consistency check on each triplet of events, where TLINK pq links events p and q, pr links events p, r, and qr links events q, r.

4 The Results

In our empirical experiments, we used an environment based on the TempEval-3 challenge (UzZaman et al., 2013). There, participants trained their classifiers on a training data set and tested them on the so called "Platinum" dataset. The Platinum dataset consists of twenty newsfeeds, which were manually annotated by experts with respect to the events, time expressions, and their temporal relations. The participants kindly provided their classifications for the the Platinum dataset to us. Table 1 summarises their precision, recall, and F1 score on the Platinum dataset, as computed using *temporal_evaluation.py* of (UzZaman, 2012). Henceforth, we denote the participants' classifiers as "individual classifiers", to distinguish them from our ensemble classifiers.

Notice that the TempEval-3 dataset has been split into two parts, rather than three, which complicates the evaluation of ensemble methods. Specifically, the individual classifiers were trained on the training set, and evaluated on the Platinum set. One can also train the ensembles on the Platinum set, but there are no other data, on which to test the ensembles. Further, it is customary to report the F1 score on the whole of the Platinum set, as used in the challenge. We hence propose two procedures for the evaluation of the ensembles, which circumvent the difficulties in two different ways.

In the first of the two procedures, the individual classifiers were trained on the training set and tested on the Platinum set to obtain their F1 scores. In Table 3, we summarise the *training* F1 score, precision, and recall of several variant ensembles on the Platinum set of 20 newsfeeds. Clearly, an ensemble (C2, C4, U4+5, N1+2) using 6 of the 11 individual classifiers (ClearTK-2, ClearTK-4, UTTime-4, UTTime-5, and Navytime-1 and Navytime-2), results in the best performance with F1 of 0.3899 and recall of 0.5. These *training* scores are considerably better than the *test* performance of the individual classifiers, as summarised in Table 1, but we do note that this comparison is not quite fair.

In the second of the two procedures, we split the Platinum dataset into two subsets, S1 and S2, of ten newsfeeds each. The individual classifiers, which were trained on the original training set, were evaluated on S1, as detailed in Table 2. We used the F1 score of the individual classifiers on S1, the first subset, as weights in the ensembles, which were tested on S2, the second subset. Table 4 details the *test* F1 score, precision, and recall of several ensembles, as tested on the 10 newsfeeds of S2. These *test* scores on S2 are considerably better than the *test* performance of the individual classifiers on $S2$, as presented in Table 2. Notice that the F1 score of each individual classifier is worse on S2 than on $S1 \cup S2$, as per Tables 1–2, which may suggest why results of the ensembles often worse on S2 than on $S1 \cup S2$, as per Tables 3–4. Overall, the results of the ensembles on $S2$ in Table 4 provide convincing evidence of the impact of the ensemble method.

Throughout the two procedures, the initial ensemble was composed of three classifiers, ClearTK-2, UTTime-4, and Navytime-1, one from each of the classifier groups. This ensemble is labelled C2, U4, N1, using the IDs introduced in Table 1. Subsequently, classifiers were added one at a time, in a complete enumeration of the combinations. Notice that the ensemble C2, U4, N1, which is diverse, outperforms C1, C2, C3, C4, where the classifiers provide similar results, albeit with higher F1 scores to start with. This has been the outcome in both procedures, supporting the assertion that diversity matters.

Both procedures were implemented in Python. The use of Pyomo, the Python-based mathematical modelling tool (Hart et al., 2012), allowed us to switch between cbc, the open-source integer programming solver (Forrest and Lougee-Heimer, 2005), and IBM ILOG CPLEX 12.4, the state-of-the-art commercial solver. Using CPLEX running on an Intel Xeon E7458 processor clocked at 2.4GHz, we were able to solve even the largest instances in seconds, as detailed in Table 5. This

Table 1: The performance of the individual classifiers on the Platinum dataset.

ID	Ensemble	F1	Precision	Recall
C1	cleartk-1	0.3517	0.3764	0.3300
C2	cleartk-2	0.3624	0.3732	0.3521
C3	cleartk-3	0.3421	0.3336	0.3510
C4	cleartk-4	0.3594	0.3526	0.3664
N1	navytime-1	0.3079	0.3519	0.2737
N2	navytime-2	0.3588	0.5078	0.2774
U1	UT-1	0.2428	0.1490	0.6556
U2	UT-2	0.2415	0.1487	0.6424
U3	UT-3	0.2422	0.1507	0.6170
U4	UT-4	0.2882	0.3752	0.2340
U5	UT-5	0.3499	0.3605	0.3400

Table 2: The performance of the individual classifiers on S2, the subset of Platinum.

ID	Classifier	F1	Precision	Recall
C1	cleartk-1	0.3140	0.3302	0.2993
C2	cleartk-2	0.3358	0.3398	0.3318
C3	cleartk-3	0.3240	0.3144	0.3341
C4	cleartk-4	0.3331	0.3320	0.3341
N1	navytime-1	0.2760	0.3145	0.2459
N2	navytime-2	0.2347	0.2899	0.1972
U4	UT-4	0.2599	0.3074	0.2251
U5	UT-5	0.3293	0.3224	0.3364

should not be surprising, considering the largest instance (WSJ_20130322_159) for 11 individual classifiers had dimension of the decision variable as low as 7365, but suggests the scalability of such an approach.

5 Conclusions

Building an ensemble of classifiers can provide a significant improvement in precision and recall over the individual classifiers. The improvement in recall is understandable, considering the ensemble uses the *union* of classifiers' results. We have demonstrated that enforcing consistency constraints over the results of even a small number of individual classifiers improves precision in this particular case. Further, this overall consistency is required in order to make practical use of the inferred temporal relations between events in real-world applications.

Promising directions for future research include exploring alternative weights such as precision, means of combining the individual classifiers, c.f. Kittler et al. (1998), the use of soft constraints, c.f. Burke et al. (2012), and applications to spatial expressions, where a similar algebra can be developed.

Table 3: The results of the first procedure with ensembles described by IDs introduced in Table 1.

IDs	F1	Precision	Recall
C1–3	0.3583	0.3515	0.3653
C1–4	0.3600	0.3488	0.3720
C2, C4, U4, N1	0.3756	0.3126	0.4702
C2, U4, N1	0.3786	0.3205	0.4625
C2, C4, U4+5, N1	0.3862	0.3182	0.4912
C1–4, U4+5, N1+2	0.3877	0.3166	0.5000
C1–2, C4, U4+5, N1+2	0.3893	0.3192	0.4989
C2, C4, U4+5, N1+2	0.3899	0.3195	0.5000

Table 4: The results of the second procedure with ensembles described by IDs introduced in Table 2.

IDs	F1	Precision	Recall
C2, U4, N1	0.3495	0.2857	0.4501
C2+4, U4, N1	0.3586	0.2941	0.4594
C2+4, U4+5, N1+2	0.3602	0.2917	0.4710
C1–4, U4+5, N1+2	0.3643	0.2952	0.4756
C2+4, U4+5, N1	0.3649	0.2988	0.4687
C1+2, C4, U4+5, N1+2	0.3671	0.2962	0.4826

Table 5: Dimensions of the integer programs and run-time of two well-known integer programming solvers, reconciling all 11 classifiers in Table 1.

Instance	Dim.	cbc (secs)	CPLEX (secs)
AP_20130322	2,790	8.01	0.57
bbc_20130322_1150	4,155	23.94	1.02
bbc_20130322_1353	4,440	119.1	3.73
bbc_20130322_1600	2,535	16.37	1.49
bbc_20130322_332	3,675	8.56	1.01
bbc_20130322_721	2,115	11.16	0.69
CNN_20130322_1003	7,200	181.49	4.69
CNN_20130322_1243	900	0.9	0.14
CNN_20130322_248	2,025	2	0.71
CNN_20130322_314	3,045	19.86	1.19
CNN_20130322_821	285	0.04	0.01
nyt_20130321_cyprus	6,285	82.55	4.11
nyt_20130321_china	4,755	55.61	3.76
nyt_20130321_sarkozy	2,130	11.16	2.43
nyt_20130321_women	4,950	41.63	4.04
nyt_20130321_strange	2,820	34.82	1.89
WSJ_20130318_731	2,025	5.76	0.43
WSJ_20130321_1145	1,965	6.58	0.64
WSJ_20130322_804	3,270	81.87	2.56
WSJ_20130322_159	7,365	n/a	75.57

References

James F Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

Steven Bethard. Cleartk-timeml: A minimalist approach to tempeval 2013. In *Second Joint Conference on Lexical and Computational Semantics (* SEM)*, volume 2, pages 10–14, 2013.

Edmund K Burke, Jakub Mareček, Andrew J Parkes, and Hana Rudová. A branch-and-cut procedure for the udine course timetabling problem. *Annals of Operations Research*, 194(1):71–87, 2012.

Nate Chambers. Navytime: Event and time ordering from raw text. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2*, pages 73–77. Association for Computational Linguistics, 2013.

Nathanael Chambers and Daniel Jurafsky. Jointly combining implicit constraints improves temporal ordering. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 698–706. Association for Computational Linguistics, 2008.

Quang Xuan Do, Wei Lu, and Dan Roth. Joint inference for event timeline construction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 677–687. Association for Computational Linguistics, 2012.

Radu Florian, Silviu Cucerzan, , Charles Schaefer, and David Yarowsky. Combining classifiers for word sense disambiguation. *Natural Language Engineering*, 8: 327–341, 12 2002.

John Forrest and Robin Lougee-Heimer. Cbc user guide. In *Tutorials in Operations Research*, pages 257–277. INFORMS, 2005.

Goran Glavaš and Jan Šnajder. Construction and evaluation of event graphs. *Natural Language Engineering*, to appear:1–46, 2014.

William E Hart, Carl Laird, Jean-Paul Watson, and David L Woodruff. *Pyomo – Optimization modeling in python*. Springer Science & Business Media, 2012.

Josef Kittler, Mohamad Hatef, Robert P.W. Duin, and Jiri Matas. On combining classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3): 226–239, 1998.

Peter B Ladkin. *Constraint reasoning with intervals: a tutorial, survey and bibliography*. International Computer Science Institute, 1990.

Natsuda Laokulrat, Makoto Miwa, Yoshimasa Tsuruoka, and Takashi Chikayama. Uttime: Temporal relation classification using deep syntactic features. In *Second Joint Conference on Lexical and Computational Semantics (* SEM)*, volume 2, pages 88–92, 2013.

Vasin Punyakanok, Dan Roth, Wen-tau Yih, and Dav Zimak. Semantic role labeling via integer linear programming inference. In *Proceedings of the 20th international conference on Computational Linguistics*, page 1346. Association for Computational Linguistics, 2004.

Dan Roth and Wen-tau Yih. A linear programming formulation for global inference in natural language tasks. In *Proceedings of CoNLL-2004*, 2004.

Roser Saurí, Lotus Goldberg, Marc Verhagen, and James Pustejovsky. Annotating events in English. TimeML annotation guidelines, 2009. Brandeis University. Version TempEval-2010.

Naushad UzZaman. *Interpreting the Temporal Aspects of*. Thesis, University of Rochester, 2012.

Naushad UzZaman, Hector Llorens, James Allen, Leon Derczynski, Marc Verhagen, and James Pustejovsky. Semeval-2013 task 1: Tempeval-3: Evaluating time expressions, events, and temporal relations. In *Second Joint Conference on Lexical and Computational Semantics (* SEM)*, pages 1–9. Association for Computational Linguistics, 2013. Also see preprint arXiv:1206.5333.

K. Woodsend and M. Lapata. *Learning to simplify sentences with quasi-synchronous grammar and integer programming*, pages 409–420. Association for Computational Linguistics, 2011.